



RESEARCH DEPARTMENT



THE QUEEN'S AWARD
TO INDUSTRY

REPORT

A digital telecine processing channel

A. Oliphant, M.A.

A DIGITAL TELECINE PROCESSING CHANNEL
A. Oliphant, M.A.

Summary

An experimental digital telecine processing channel is described. It performs the operations of gamma correction, colour masking and aperture correction. To give sufficient coding resolution in the green channel a dual-range ADC has been developed which gives 11-bit resolution for small signals.

The problem of moiré patterns arising from aliasing caused by non-linear processing have been investigated and found to be insignificant.

A subjective test has shown that 10 bits per sample are needed in the green channel; with this resolution, dither is unnecessary.

Issued under the authority of



Research Department, Engineering Division,
BRITISH BROADCASTING CORPORATION

Head of Research Department

A DIGITAL TELECINE PROCESSING CHANNEL

Section	Title	Page
	Summary	Title Page
1.	Introduction	1
2.	Requirements of a digital channel	1
	2.1. Component units	1
	2.2. Number of bits needed	4
	2.3. Sampling frequency	5
3.	Design of a digital channel	6
	3.1. The range-changing ADC	6
	3.2. Logarithmic converter	7
	3.3. Three-product adder	7
	3.4. Exponential converter	9
4.	Subjective tests	9
	4.1. Equipment	9
	4.2. Test pictures	10
	4.3. Test procedure	10
5.	Results of subjective tests	11
6.	Conclusions	11
7.	References	11

A DIGITAL TELECINE PROCESSING CHANNEL

A. Oliphant, M.A.

1. Introduction

A telecine signal processing channel performs a number of operations on the signals derived from the film. In a flying spot equipment these are:

1. Afterglow correction — to correct for the afterglow of the phosphor in the flying spot tube.
2. Aperture correction — to increase the high frequency content of the signal so as to compensate for various effects which spread the image of the scanning spot, thus softening the displayed picture.
3. Low-pass filtering — to remove out-of-band components.
4. Masking — to correct for the errors in colour rendering introduced by the film and the telecine analysis.
5. Gamma correction — to precorrect for the non-linearity of the display tube.
6. Blanking insertion.

A block diagram of an analogue processing channel is shown in Fig. 1.

Masking is carried out on logarithmic signals. This is because masking is used mainly to correct for errors in density of the three dye layers in the film. The signal from a photocell is proportional to the transmission (T) of the film, which is related to density by the formula $D = -\log(1/T)$.

Gamma correction is a power-law process, so it is most easily carried out by adjusting the gain of the signal path between the logarithmic and exponential amplifiers.

Most of the above processing could be done digitally, with the following advantages:

1. Easier manufacture and commissioning: the large number of preset adjustments needed in analogue channels could be greatly reduced.
2. Improved stability, reducing routine maintenance.
3. Improved reliability.
4. Improved accuracy of processing, with less variation from channel to channel.
5. Reduced noise and spurious signals.
6. Easier interfacing with other digital equipment; for example, a noise reducer.

7. Possible future cost savings as digital integrated circuits become cheaper.

Digital afterglow correction does not appear at present to offer many advantages. Analogue afterglow correctors are simple and stable, whereas a digital corrector would involve an extremely complicated digital filter. Such a filter would be expensive and also difficult to adjust for the afterglow characteristics of different flying spot tubes and the effects of tube ageing.

Although the channel was built to process telecine signals, most of the units described could also be used in a camera processing channel.

2. Requirements of a digital channel

2.1. Component units

A digital telecine channel must perform the same functions as an analogue channel. But for easier processing, the operations may be carried out in a different order or in a different way. A block diagram of a digital channel is shown in Fig. 2. The output from each afterglow corrector is converted to digital form in an analogue-to-digital converter (ADC). Each converter contains a low-pass filter to remove components above 5.5 MHz; if these were not removed they would cause aliasing in a sampled system. The specification of the ADC depends on which colour channel it is used in, that for the green channel being subject to the most stringent requirements because errors in this channel are known to be the most visible.¹

A logarithmic converter follows the ADC in each channel. This converts linearly coded signals to logarithmic form for masking and gamma correction. In concept, masking and gamma correction involve a matrix multiplication which may be expressed as follows:

$$\begin{pmatrix} \log R_{out} \\ \log G_{out} \\ \log B_{out} \end{pmatrix} = \begin{pmatrix} k_{11} & k_{12} & k_{13} \\ k_{21} & k_{22} & k_{23} \\ k_{31} & k_{32} & k_{33} \end{pmatrix} \times \begin{pmatrix} \log R_{in} \\ \log G_{in} \\ \log B_{in} \end{pmatrix}$$

In practice, a three-product adder in each channel multiplies each of the three logarithmic colour separation signals by appropriate constants and adds the products together.

Exponential converters turn each resulting signal into a gamma-corrected masked, colour signal.

The digital aperture corrector used in the channel is described in another Report.² It generates a luminance correction signal which is added to all three colour signals after the exponential converters. The adder stages also clip and blank the signals.

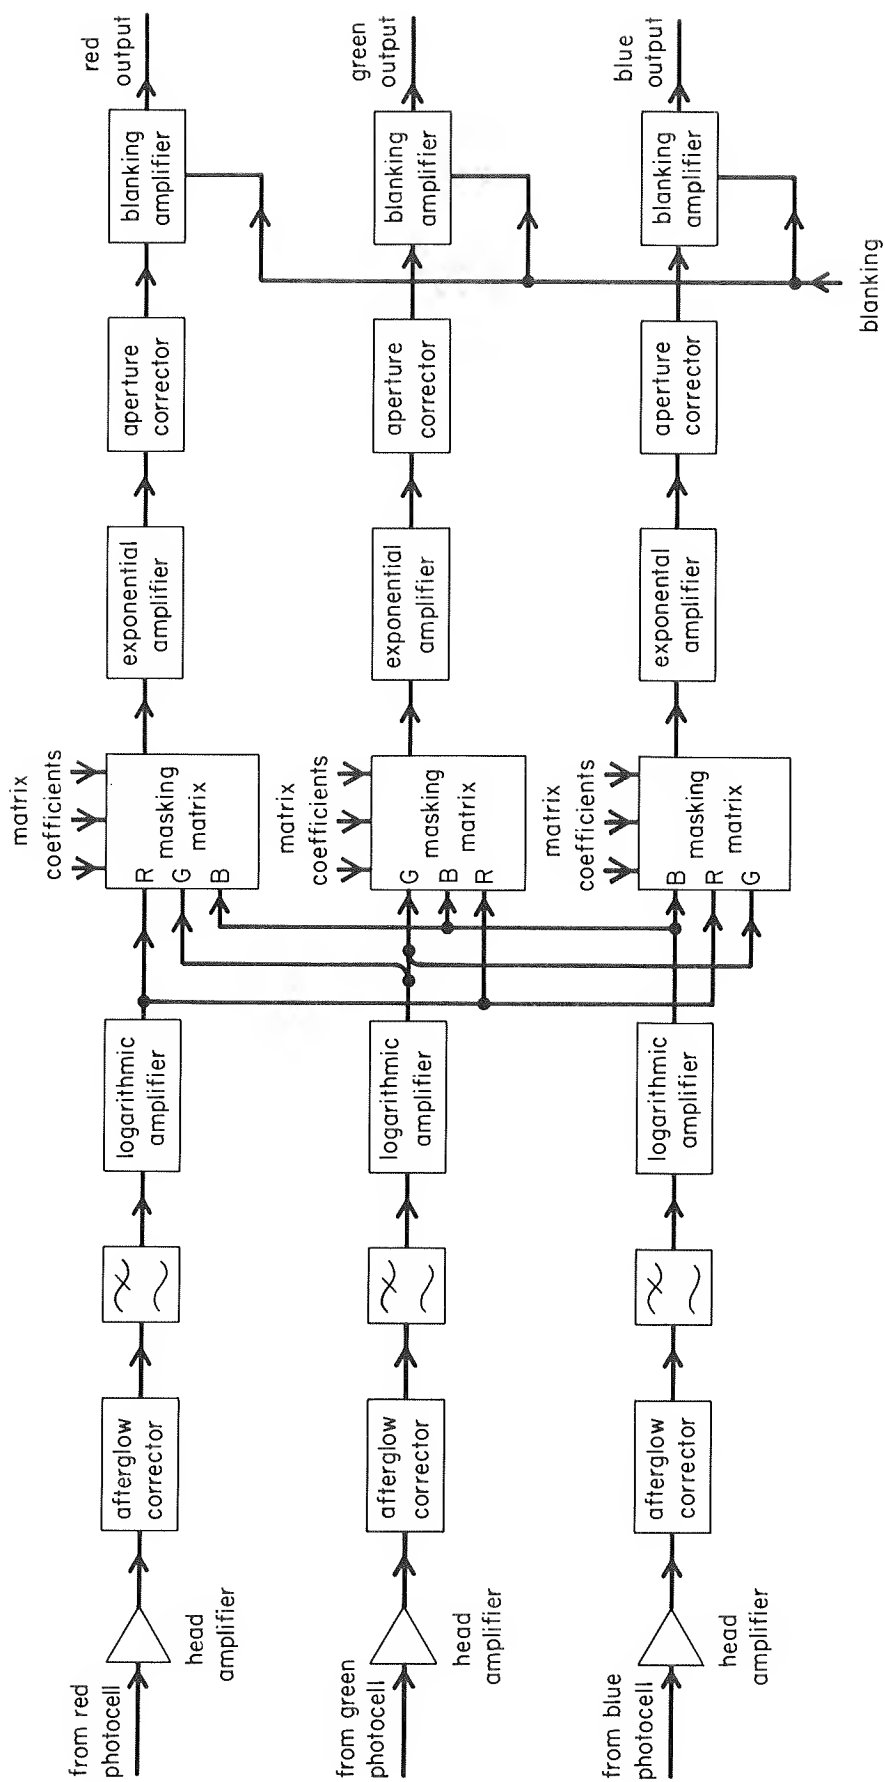


Fig. 1 - Block diagram of an analogue processing channel





Fig. 2 - Block diagram of a digital processing channel

Fig. 3 - Visibility of quantising after digital gamma correction



2.2. Number of bits needed

In most digital video processing eight bits are used to represent each sample; this number has been found to be enough to reduce the effects of quantising to an acceptable level when the signals being digitally coded are composite and have already been gamma corrected.

When signals are digitally coded before gamma correction, the visibility of quantising is increased in lowlights: this is because gamma correction increases the gain for signals near black. Low-level signals therefore need to be coded to more than the usual eight bits coding resolution: * previous work¹ suggests that ten bits per sample might be needed, at least in the green channel, even allowing for the extra bit gained by coding colour separation rather than composite signals. As quantising errors are not as visible in the red and blue channels, eight bits per sample, together with dither, would probably suffice.

Video ADCs available at present code signals to a resolution of only eight bits per sample; this presents a problem in providing the extra resolution needed in the green channel. However, the extra resolution is needed only when the input signal is small. To achieve greater resolution without the greatly increased complexity of developing an ADC of at least 10-bit resolution, an eight bit ADC³ has been modified so as to insert analogue preamplification whenever the input signal falls below a pre-determined threshold. The preamplification factor is an exact power of 2; thus the 8-bit word from the ADC can be located within a longer word by simply displacing it by the appropriate number of binary places.

* The resolution of a digital coding system is the relative value of the least significant bit (l.s.b.). It may be quoted as '1 part in 2^n ' or simply as ' n bits'. The resolution is distinguished from the accuracy which is the difference between the actual and ideal analogue signal values corresponding to a given digital signal, compared with the full scale signal value. Accuracy is generally quoted as plus or minus a fraction of a l.s.b.

Fig. 3 shows the visibility of quantising effects on signals that have been linearly coded and subsequently gamma corrected; the figure is a re-drawn version of Fig. 8 of Ref. 1 and relates to eight-bit coding and a gamma of 0.45. The visibility is shown as a fractional change in perceived luminance, $\Delta Y/Y$.

The smallest change in luminance that can be seen is about 2% (the Fechner fraction). Fig. 3 shows that $\Delta Y/Y$ is greater than 2% when the input signal falls below 53% of peak level, but does not start to rise rapidly until the signal level falls below about 15%. A preamplification factor of 8 gives three extra bits at signal levels below 12.5%; this reduces $\Delta Y/Y$ so that its maximum value in the preamplified range is about 2% and the maximum value elsewhere is about 4%. It would be possible to reduce this value of 4% by introducing an intermediate preamplification factor of 2 which would be used at signal levels between 12.5% and 50%. However, this extra complication would not be justified since the visual improvement would be so small that it would not be noticeable.

As mentioned above, eight bits per sample plus dither would probably be enough to make quantising effects invisible in the red and blue channels. But this statement ignores the effects of masking and aperture correction, which cause quantising effects in the red and blue channels to appear at lower levels in the green output signals. When a medium-saturation mask is used, as in the subjective tests described in Section 4, this cross-contouring effect is not significant. However, it may be necessary to use more than eight bits per sample in the red and blue channels if a large degree of masking, or considerable aperture correction, is required.

Arithmetic processing in the channel after coding will give rise to extra bits, which cannot all be thrown away without increasing the quantising noise. Twelve bits are used in the processing units to ensure eight-bit accuracy at the output.

2.3. Sampling frequency

The sampling frequency (f_s) used in a digital system normally must be greater than twice the highest frequency present in the analogue signal (the Nyquist limit). If this criterion is not met, frequencies greater than half the sampling frequency ($f_s/2$) will be folded back (or 'aliased') about $f_s/2$. The alias components may beat with high frequencies in the signal to produce lower frequency moiré patterns in the picture.

It is often convenient in digital television systems to use a sampling frequency which is an integral multiple of the colour subcarrier frequency. The lowest multiple which is above the Nyquist limit is three times subcarrier (about 13.3 MHz), and the digital channel described in this Report was designed to work at this frequency or alternatively at 851 times line frequency, (also about 13.3 MHz). It is, however, possible to exploit the structure of the spectrum of the television signal to allow sampling at twice colour subcarrier frequency (about 8.86 MHz).⁴ In the future this frequency may be used for the transmission of digital signals; if so it might then be convenient to sample the signal in the processing channel at four times subcarrier frequency; however, at four times subcarrier frequency (about 17.7 MHz) the design of the digital processing logic would become much more difficult, using devices available at present.

Another factor which could influence the choice of sampling frequency is the production of alias components by non-linear processing. Any non-linear process generates harmonics of the input signal. If these harmonics occur at frequencies above ($f_s/2$) they may cause moiré patterns due to aliasing, as explained above.

Gamma correction is a non-linear process. Fig. 4 illustrates (in the time domain) how moiré patterns are produced in a digital gamma corrector. Fig. 4(a) shows a 100% sine wave signal of $1/3$ sampling frequency sampled in two different phases; Fig. 4(b) shows that the mean level of the gamma-corrected output samples depends on the sampling phase. If the signal frequency is not $1/3$ sampling frequency, the sampling phase will vary periodically. As the phase changes, the mean level will change, producing a low-frequency beat added to the high-frequency signal. Similar beats are, in principle, produced by inputs having frequencies near all submultiples of sampling frequency.

The amplitude of the beat depends on the input signal frequency, its amplitude, and the sit (added d.c.). The table below shows how the amplitude of the beat varies with input frequency for a given value of sampling frequency (f_s). The figures are for a 100% sine wave signal and a gamma of 0.5.

Input frequency	$f_s/2$	$f_s/3$	$f_s/4$	$f_s/5$	$f_s/6$
Beat level (dB)	-13	-21	-26	-30	-33




Fig. 4(a) - 100% sine wave inputs, sampled at different phases
Fig. 4(b) - Gamma corrected output showing variation in mean level of samples

The amplitude of the beat also decreases very rapidly as the lowest part of the sine wave departs from black level – that is, as the sit of the signal is increased.

Clipping is another non-linear process which generates harmonics of the input signal. The amplitude of the harmonics generated by clipping will be greater than those generated by gamma correction, which is a smooth non-linearity generating significant harmonics only with large input signals. It should be noted that clippers in analogue channels can generate moiré patterns on vertical detail because of the sampling action of the 625 line scanning system, which, on a field-by-field basis, is spatially equivalent to a sampling frequency of 7.4 MHz. This moiré from vertical detail is not usually considered to be a serious impairment.

Tests with the digital channel using sampling at either three times subcarrier frequency or 851 times line frequency have shown that moiré effects due to gamma correction are just visible with electronically generated sine wave signals, but very rarely visible with real pictures; moiré from clipping is easily visible with electronically generated signals, but can be seen with real pictures only if the signals are grossly distorted by severe clipping. Thus the generation of beats in a non-linear digital processing channel and the corresponding appearance of moiré patterns is not a problem if a sampling frequency of about three times subcarrier frequency, or higher, is used.

3. Design of a digital channel

3.1. The range-changing ADC*

As explained in Section 2.2, an eight bit video ADC³ was modified to give eleven-bit resolution when coding small signals. The three extra bits were generated by amplifying the analogue signal by a factor of eight (2^3) when the signal amplitude was less than about one eighth of its maximum value; this changed the effective coding range of the ADC. The eight bit output of the ADC was shifted by three binary places when the preamplification was switched in, thus giving an eleven bit output.

The circuit associated with the modification was designed as a replacement for the existing sample-and-hold board of the ADC. A block diagram of the new sample-and-hold board is shown in Fig. 5. The new board consists of two separate paths, one with eight times the gain of the other, each path containing a diode bridge to sample the signal. A comparator switches the sampling pulses from one bridge to the other as the input video signal level passes a preset threshold level. Since each sampling bridge appears as a high impedance when no sampling pulses are applied to it, the outputs of the two bridges can both be connected to the same 'hold' capacitor.

It is important that the delays through both video paths are exactly the same and that the sampling pulses

* The range-changing ADC was designed by R.P. Marsden.

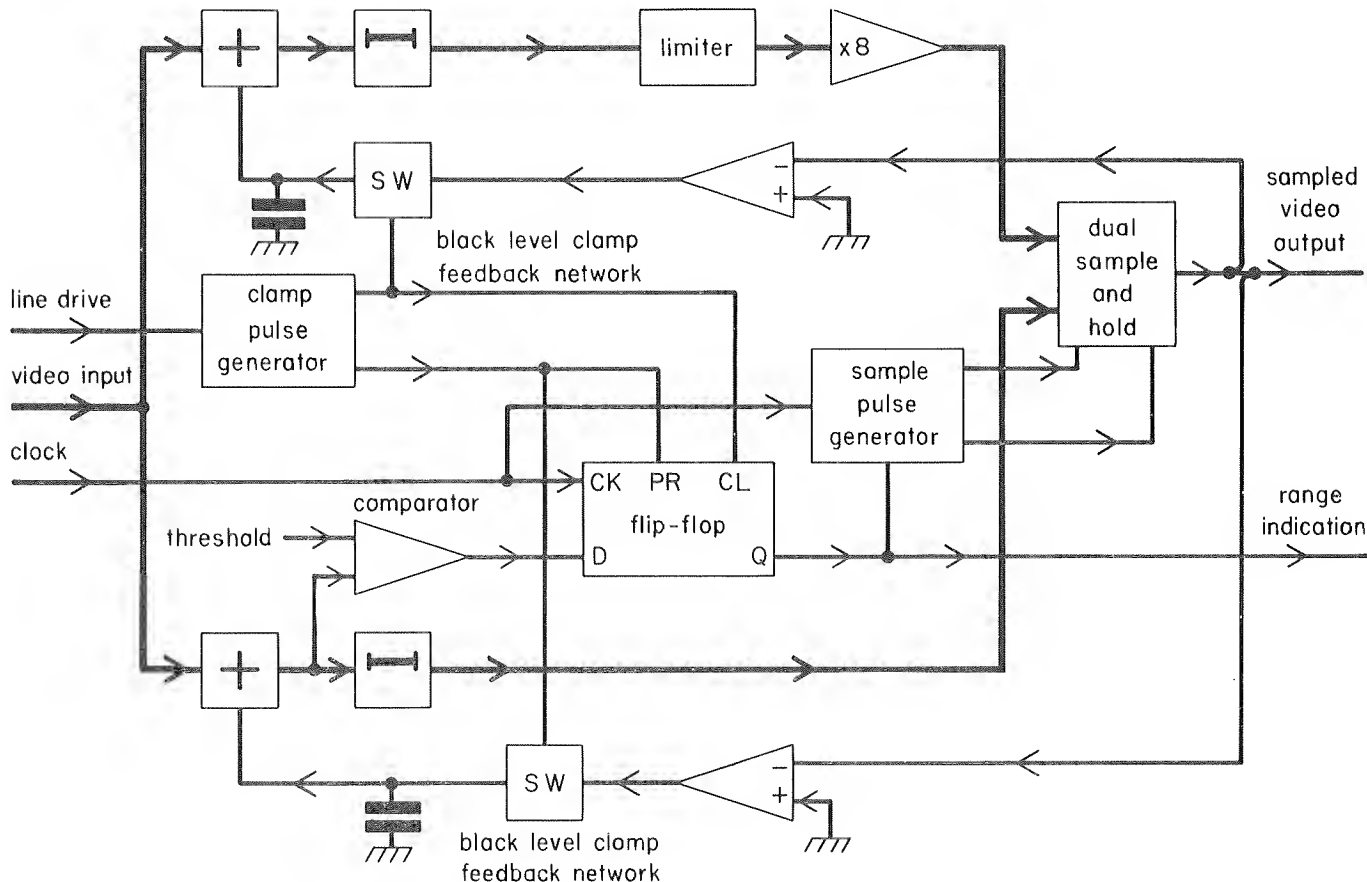


Fig. 5 - Modified sample-and-hold

change from one bridge to the other at the same point, in relation to the video waveforms presented to both bridges, as the comparator makes the decision to switch. The delay through the switching logic is much greater than that through the video paths, so equalising delay lines are needed in each video path.

Each video path has a feedback clamp which encompasses the sample-and-hold stage. The clamp works by sampling the level of the signal at the output of the dual sample-and-hold during the back porch; the samples are compared with a reference level and the difference is used to control the d.c. level in the buffer amplifier at the input to each video path. The signal must be sampled twice in each line, once for each video path, and the sample-and-hold forced to sample the appropriate signal at that time regardless of the input signal level. This is arranged by feeding clamp pulses to the preset and clear inputs of the flip-flop that registers the state of the comparator.

The comparator threshold is set slightly below 12.5% of white level to allow for marginal decisions influenced by noise and component drift: if the high gain channel is sampled when the input signal exceeds 12.5%, the ADC will overload. Apart from this restriction, the threshold is not critical. A limiter prevents gross overloads in the later stages of the high gain channel.

The ratio of the gains of the two paths must be correct to ± 1 part in 64. An error in the ratio of the gains will change the size of the step between the highest level of the high gain (low signal) range and the lowest level of the normal range. This step must not be larger than the step above it or it will be more visible; nor must it be negative, or the steps will be in the wrong order. The height of the normal quantising steps is $1/256$ of peak level, and the threshold is at about $1/8$ of peak level. So the relative gain may vary by $8/256$ (1 part in 32) or ± 1 part in 64.

3.2. Logarithmic converter

For convenience in digital processing, logarithmic equivalents to base 2 of the samples are derived; these can be derived accurately by using Read Only Memories (ROMs). A method which needs only a small memory is to find the characteristic and mantissa of the number separately.* A block diagram of this arrangement is shown in Fig. 6. The characteristic of the logarithm is found first by detecting the position of the most significant 1 using a priority encoder. The four-bit output of the priority encoder is the characteristic of the logarithm, and is also used to shift the input number so that it appears as a number between binary 1.0000000 and 1.1111111. This number is then applied to a ROM that is programmed with the binary values of the logarithms of numbers in that range. The output of the ROM is the eight-bit mantissa of the logarithm. The complete logarithm is thus a twelve-bit number.

* If $\log_2 x = (c + m)$ where c is an integer and $0 \leq m < 1$ then c is the characteristic of the logarithm and m is the mantissa.

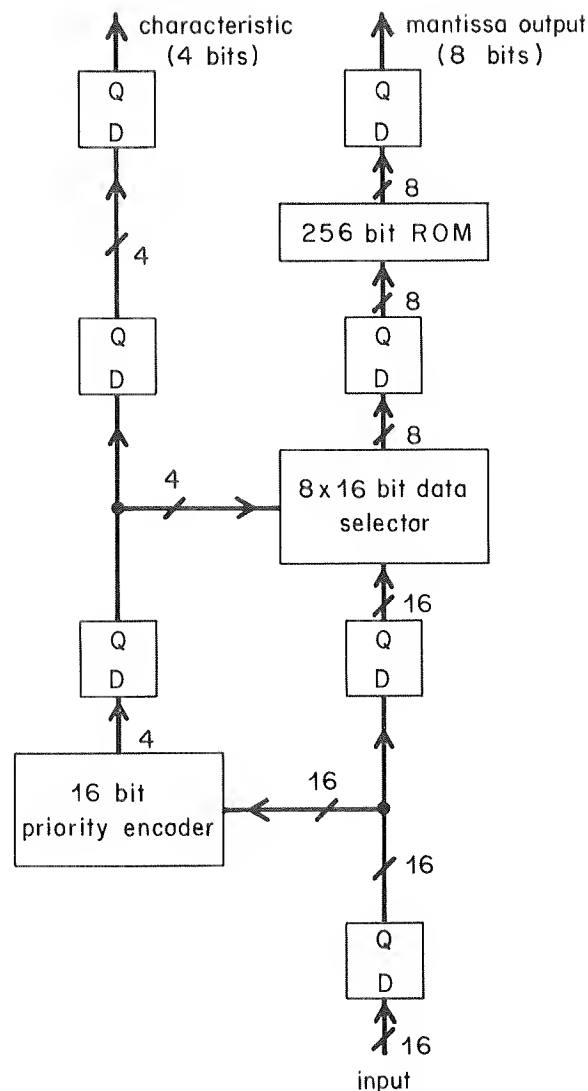


Fig. 6 - Block diagram of log converter

A problem arises if the input signal is allowed to go to zero, since the logarithm would then be $-\infty$. In practice this is overcome by adding a small offset to the input signal so that its value can never become zero.

3.3. Three-product adder*

As explained in Section 2.1, the three-product adders (one in each channel) perform the operations of colour masking and gamma correction. Each three-product adder takes in the logarithmic equivalents of the three colour separation signals, multiplies each by a constant and adds the three products together:

$$Y = k_1 x_1 + k_2 x_2 + k_3 x_3$$

The constants (the multipliers) are specified to six bits and the input numbers (the multiplicands) to twelve bits. The output is a twelve-bit number but it is only accurate to ten bits because of truncation errors in the arithmetic. The constants can be varied in the ranges

* The three-product adder was designed by I. Childs.

$$\begin{aligned}
0 \leq k_1 &\leq 1\frac{31}{32} \\
-1 \leq k_2 &\leq \frac{31}{32} \\
-1 \leq k_3 &\leq \frac{31}{32}
\end{aligned}$$

To avoid negative numbers, k_2 and k_3 are presented in offset binary, i.e. the input 000000 represents -1 . 100000 represents 0 and 111111 represents $+1$. The use of this format requires that the number $(x_2 + x_3)$ be subtracted from the final output to remove the offset.

The multiplication is performed by gating each of the multiplicands with each bit of the appropriate multiplier; this generates eighteen twelve-bit partial products of varying levels of significance. These partial products are then added together in a tree network. The process is shown diagrammatically in Fig. 7.

A three-product adder was designed and built using 27 Arithmetic Logic Units (ALUs) to perform the input gating and the first stage of addition, and 30 four-bit adders

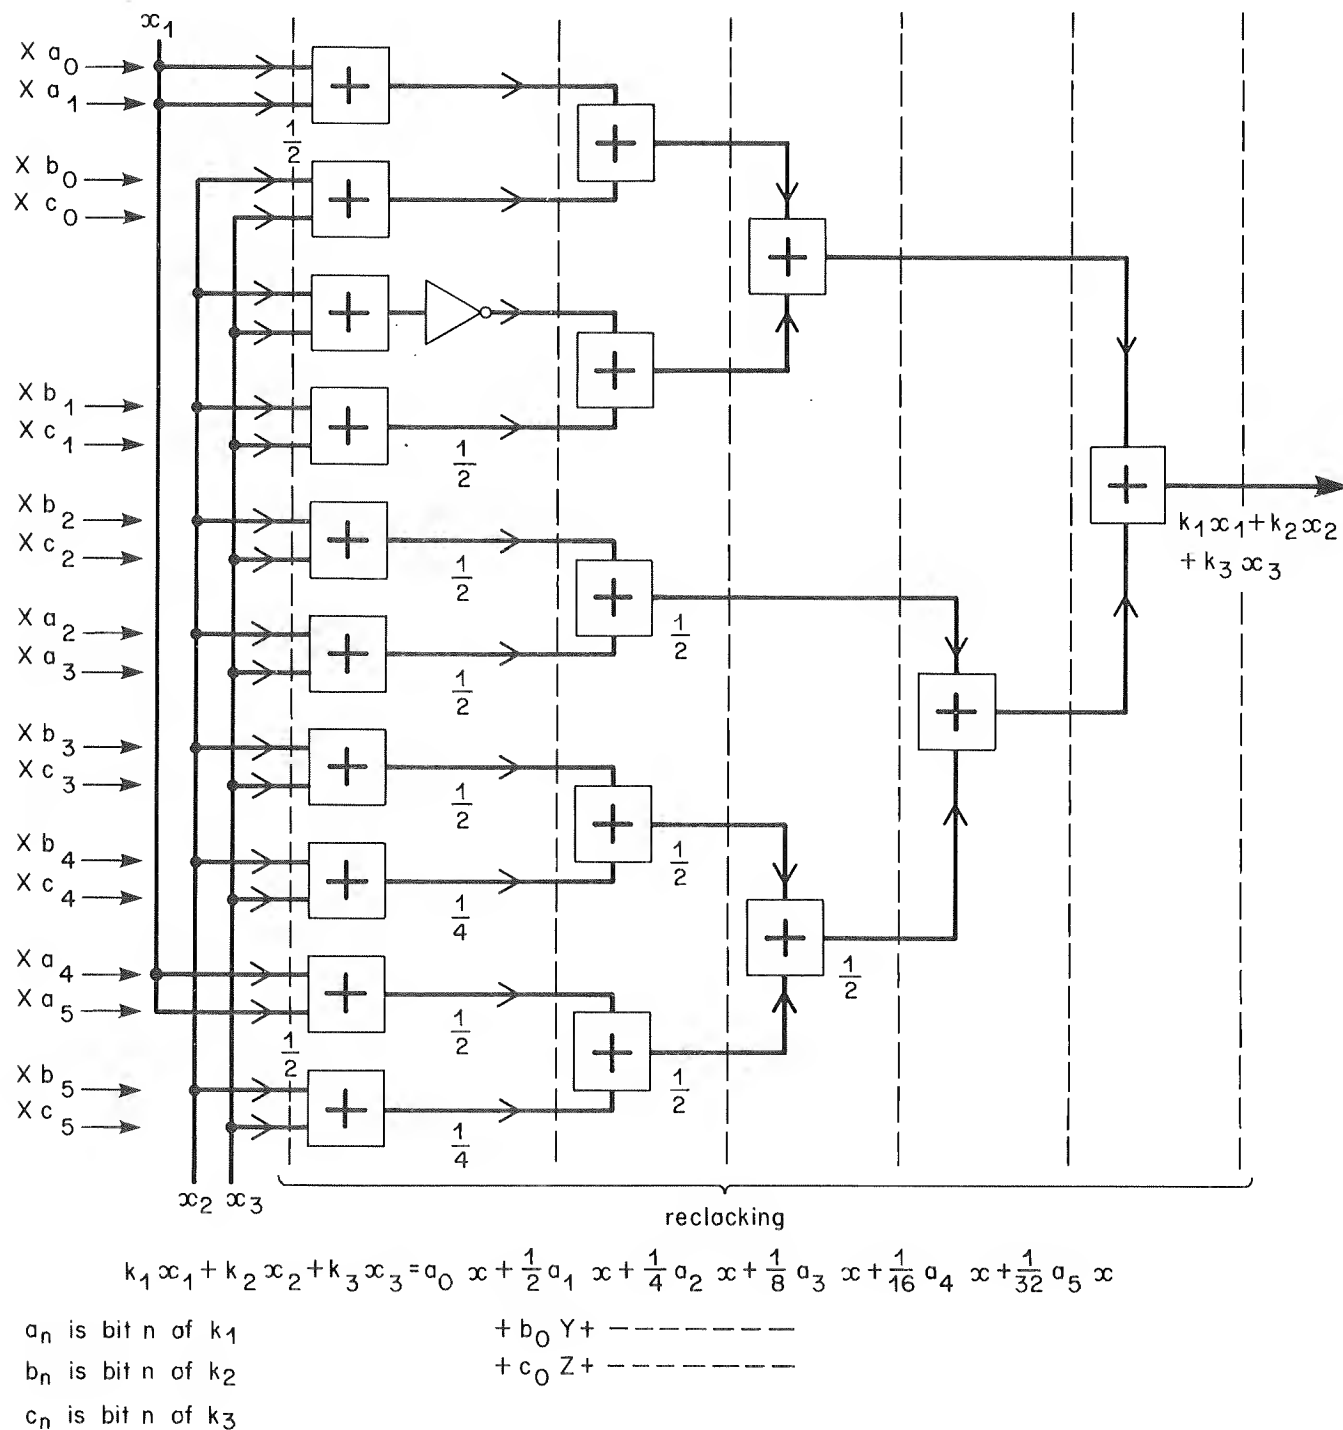


Fig. 7 - Block diagram of the three-product adder

to form the rest of the tree. For a sampling frequency of $3f_{sc}$ ($3 \times$ PAL subcarrier frequency) the data is clocked at 75ns intervals; this does not allow time to do all the processing in one clock period. The addition of the partial products is therefore carried out in stages, with storage between stages. The input store uses three 12-bit stores (in practice six 6-bit stores), and the intermediate storage between stages uses nineteen 12-bit stores (thirty-eight 6-bit stores). Thus the three-product adder is an elaborate unit: the 3×3 matrix using three such adders is expensive compared with its analogue equivalent.

3.4. Exponential converter*

The exponential converter performs the exact inverse operation to the logarithmic converter. A block diagram is shown in Fig. 8. Each eight-bit mantissa is applied to a 256-bit ROM which provides a code between 1.0000000 and 1.1111111. This is shifted to the appropriate degree of significance by a 12×8 -bit data selector controlled by the characteristic input, and forms the gamma-corrected output.

4. Subjective tests

A subjective test was performed to test the prediction from earlier work that nine or ten bits would be needed in the green channel of a digital processing channel.¹

* The exponential converter was designed by D.T. Wright.

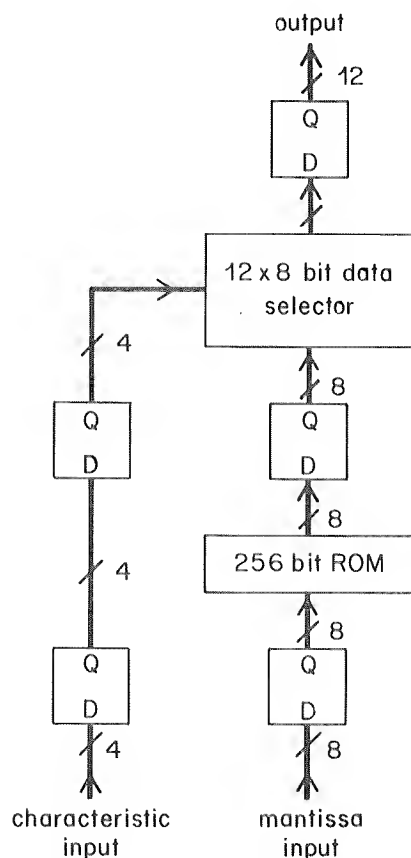


Fig. 8 - Block diagram of exponential converter

4.1. Equipment

A flying-spot slide scanner provided the linear analogue video signals for the tests. The red and blue signals were each coded to eight bits; half sampling frequency and random noise dithers were added to both signals to reduce the visibility of quantising.


The range-changing 11-bit ADC was used in the green channel. The 11-bit output from this ADC was fed to the channel input via a switching unit which allowed bits to be discarded. The switching logic was arranged so that setting a particular bit to zero forced the bit below it to a 1, thus minimising the d.c. shift caused by truncation.

The digital channel performed the operations of masking, gamma correction and aperture correction. The matrix coefficients used, representing a medium-saturation mask and a gamma of 0.4, were:

$$\begin{pmatrix} \log R_{out} \\ \log G_{out} \\ \log B_{out} \end{pmatrix} = \begin{pmatrix} 14/32 & -1/32 & 0 \\ -1/32 & 15/32 & -1/32 \\ 0 & -2/32 & 15/32 \end{pmatrix} \times \begin{pmatrix} \log R_{in} \\ \log G_{in} \\ \log B_{in} \end{pmatrix}$$



(a)



(b)

Fig. 9 - Monochrome versions of slides used in tests

(a) Boy (b) Fish-hook

The signals from the three DACs at the outputs of the digital channel were displayed, without PAL coding, on a high quality 22 in. shadowmask monitor. The peak brightness of the monitor was set to 70 cd/m² as recommended by the CCIR⁵ and its cutoff was set using the standard line-up waveform (PLUGE). The background illumination in the test room was low, and consisted mainly of diffuse daylight.

4.2. Test pictures

Two colour slides, shown in monochrome in Fig. 9, were used in the tests. These slides had been found to be rather critical in previous tests.¹ The first slide ('Boy') has a dark plain background which readily shows quantising effects in lowlights; the second slide ('Fish-hook') contains low-level high-frequency detail which is lost when the number of bits is too few.

The use of stationary pictures made the tests easily repeatable; it also gave observers time to find and examine the most critical parts of the picture. But quantising effects are more visible when there are slow brightness changes. Therefore a 1 Hz sinusoidally varying sit of amplitude 1/1000 of peak white was added to the green channel. This sit variation was virtually imperceptible when all 11 bits were used; preliminary tests showed that its addition made no difference to the assessment of pictures derived by analogue processing, but that quantising effects were increased in visibility when too few bits were used in the digital coding.

The input signal gain was set so that the full ADC conversion range was used to within 1 or 2 dB.

4.3. Test procedure

A total of 10 technical observers took part in the tests. They were seated at distances of between five and seven times picture height from the screen.

Each session consisted of 40 tests occupying about 30 minutes in all. Each of the ten test conditions (using 7, 8, 9, 10 or 11 bits in the green channel, each with and without dither) was shown twice for each slide. The two slides were shown alternately, but the test conditions were presented in random order. The picture produced by each test condition was shown for 20 seconds and the screen was blacked out for 10 seconds between tests.

Before each session the observers were shown both an unimpaired picture and one including the impairments caused by quantisation and dither using the worst test condition (7 bits). For the assessments, the observers were asked to grade the degree of picture impairment produced by each test condition, using the CCIR 5-point impairment scale.

Grade	Impairment
5	Imperceptible
4	Perceptible, but not annoying
3	Slightly annoying
2	Annoying
1	Very annoying

Fig. 10 - Subjective test results





TABLE 1

Results of Subjective Tests

Test No.	Slide	No. of bits	Dither	Nos. in test sequence	Mean grade (10 obs)	Standard deviation
1	(a) Boy	7	OFF	15, 31	1	0
2		8		5, 39	1.55	0.69
3		9		23, 27	3.3	0.83
4		10		1, 37	4.6	0.50
5		11		9, 21	4.3	0.59
6	(a) Boy	7	ON	13, 29	1.63	0.65
7		8		3, 17	3.38	1.01
8		9		11, 25	4.23	0.60
9		10		7, 19	4.38	0.78
10		11		33, 35	4.53	0.52
11	(b) Fish-hook	7	OFF	26, 36	3.0	1.27
12		8		20, 34	2.85	1.03
13		9		4, 32	3.58	1.14
14		10		18, 40	4.73	0.44
15		11		2, 16	4.75	0.44
16	(b) Fish-hook	7	ON	10, 24	2.3	0.75
17		8		8, 12	3.0	0.84
18		9		14, 38	4.23	0.64
19		10		22, 30	4.33	0.65
20		11		6, 28	4.53	0.79

5. Results of subjective tests

The results of the subjective tests are set out in Table 1 and plotted in Fig. 10.

They show that observers found little difference between pictures produced using 10 or 11 bits in the green channel. Dither made no difference under these circumstances, and 9 bits with dither was almost as good as 10 or 11 bits.

One anomaly deserves comment. Most observers found the pictures from slide 'b' (Fish-hook) much less annoying than those from slide 'a' (Boy), when using seven bits without dither in the green channel. Indeed, at seven bits 'Fish-hook' without dither was found less annoying than with dither. This is because most observers failed to notice the loss of low-level detail, although this was pointed out before the tests as an effect of coarse quantisation. In contrast, the picture from slide 'a' (Boy) showed severe contouring across the background when using seven bits without dither in the green channel. All observers found this very annoying.

The probable reason why dither made no difference to the 10 and 11 bit signals is that the noise accompanying the signals from the scanner (mostly from the head amplifiers) was sufficiently high in level to conceal the differences between 10 and 11 bit quantisation effects. Previous work⁶ has shown that quantising will be invisible if the video signal-to-noise ratio is worse than $(6n + 5)$ dB, where

n is the number of bits used. The r.m.s. noise level from the head amplifiers is about -65 dB (rel. 0.7V), so contouring should be seen in lowlights only when $6n + 5 < 65$, i.e. when $n < 10$, when no dither is used. Thus one would expect that dither would improve 9 bit pictures, but would make little difference to 10 or 11 bit pictures.

6. Conclusions

A digital telecine processing channel has been built to perform the operations of gamma correction, masking and aperture correction. To give improved coding resolution in the green channel, a range-changing ADC was developed which gives 11 bit resolution for low-level signals.

The channel has been used to confirm the prediction from previous work that 10 bits are needed for the green channel if the signals are digitally coded before gamma correction. A subjective test has shown that nine bits with dither would be just sufficient but that ten bits should give a margin of safety. If ten bits are used, dither is probably unnecessary. The generation of moiré patterns by alias components resulting from non-linear processing has been found to be insignificant provided sampling frequencies greater than about 13.3 MHz are used.

7. References

1. WESTON, M. 1975. Pulse code modulation of video signals: visibility of level quantising effects in process-

- ing channels. BBC Research Department Report No. 1975/31.
2. WESTON, M. 1977. Digital aperture correction. BBC Research Department Report No.1978/10
 3. DEVEREUX, V.G. 1970. Pulse code modulation of video signals: 8 bit coder and decoder. BBC Research Department Report No. 1970/25.
 4. DEVEREUX, V.G. *and* PHILLIPS, G.J. 1974. Bit rate reduction of digital video signals using differential p.c.m. techniques. IBC 1974. IEE Conference Publication No. 119, pp. 83 – 87.
 5. CCIR 1974. Method for the subjective assessment of the quality of television pictures. Recommendation 500.
 6. DEVEREUX, V.G. Pulse code modulation of video signals: subjective study of coding parameters. BBC Research Department Report No. 1971/40.